

# Claims

- [c1] 1. In a database system, a method for recommending database indexes to be created for optimizing system performance, the method comprising:  
capturing a workload representative of database queries employed during system use;  
creating virtual indexes for optimizing system performance during execution of the database queries captured in the workload;  
computing cost benefits for different combinations of the virtual indexes; and  
recommending physical indexes to be created based on virtual indexes that have favorable cost benefits for the captured workload.
- [c2] 2. The method of claim 1, wherein said virtual indexes comprise in-memory data structures corresponding to sets of potential physical indexes.
- [c3] 3. The method of claim 1, wherein the capturing step includes:  
displaying a screen input button that a user may invoke to record a usage session as a workload.

- [c4] 4. The method of claim 1, wherein the workload represents user execution of a database application with a typical workload that is contemplated for the application.
- [c5] 5. The method of claim 1, wherein the workload includes information recording text of all the queries operating during the capture of the workload.
- [c6] 6. The method of claim 1, wherein the workload includes information recording settings for certain options that affect how queries are optimized.
- [c7] 7. The method of claim 1, wherein the capturing step includes:  
capturing information about a set of workloads to define a problem instance.
- [c8] 8. The method of claim 1, further comprising:  
setting a limit on how much disk space is available for physical indexes.
- [c9] 9. The method of claim 8, wherein the recommending step takes into account the limit on disk space available for physical indexes.
- [c10] 10. The method of claim 8, wherein the recommending step includes:  
if the physical indexes to be recommended for creation

exceed the limit on disk space available for physical indexes, removing some of the physical indexes from consideration.

- [c11] 11. The method of claim 10, wherein the physical indexes removed from consideration are ones having less favorable cost benefits for the captured workload.
- [c12] 12. The method of claim 10, wherein the physical indexes removed from consideration comprise at least 20 percent of bottom performing indexes considered for recommendation.
- [c13] 13. The method of claim 1, further comprising:  
specifying whether certain types of indexes should be considered at all.
- [c14] 14. The method of claim 1, wherein the creating virtual indexes step includes:  
searching for relevant indexes that will help the system's optimizer use sargable predicates for partial index scans.
- [c15] 15. The method of claim 14, wherein the optimizer creates virtual indexes without specifying ordering of columns used in sargable equality predicates.
- [c16] 16. The method of claim 1, wherein the creating virtual

indexes step includes:

searching for relevant indexes that will help provide useful orderings.

[c17] 17. The method of claim 16, wherein columns of virtual indexes may be order-independent "don't care" columns that satisfy some interesting ordering wish list of the system's optimizer.

[c18] 18. The method of claim 16, wherein columns of virtual indexes may have an unspecified sortedness.

[c19] 19. The method of claim 1, further comprising:  
collapsing some of the virtual indexes together, if feasible for the workload.

[c20] 20. The method of claim 19, wherein the collapsing step includes:

identifying that columns of one virtual index are a superset of another the columns of another virtual index, and that both indexes may be combined into a single virtual index that is feasible for the workload; and  
identifying that sortedness of a column of a virtual index, if unspecified, may be specified to allow it to be combined with an index with identical columns but specified sortedness; and  
identifying that a virtual index that has columns of op-

posite sortedness of a second virtual index, and that both indexes may be combined into a single virtual index.

- [c21] 21. The method of claim 1, further comprising:  
polling periodically in the method to ensure that the system is working with accurate cost information.
- [c22] 22. A computer-readable medium having processor-executable instructions for performing the method of claim 1.
- [c23] 23. A downloadable set of processor-executable instructions for performing the method of claim 1.
- [c24] 24. A system that recommends database indexes to be created for optimizing system performance, the system comprising:  
a database system that executes database queries; and  
an index consultant for capturing a workload representative of database queries executed during typical system use; creating virtual indexes for optimizing system performance during execution of the database queries captured in the workload; computing cost benefits for different combinations of the virtual indexes; and recommending physical indexes to be created based on virtual indexes that have favorable cost benefits for the

captured workload.

- [c25] 25. The system of claim 24, wherein said virtual indexes comprise in-memory data structures corresponding to sets of potential physical indexes.
- [c26] 26. The system of claim 24, wherein the index consultant displays a screen input button that a user may invoke to record a usage session as a workload.
- [c27] 27. The system of claim 24, wherein the workload represents user execution of a database application with a typical workload that is contemplated for the application.
- [c28] 28. The system of claim 24, wherein the workload includes information recording text of all the queries operating during the capture of the workload.
- [c29] 29. The system of claim 24, wherein the workload includes information recording settings for certain options that affect how queries are optimized.
- [c30] 30. The system of claim 24, wherein the index consultant captures information about a set of workloads to define a problem instance.
- [c31] 31. The system of claim 24, wherein the index consultant may receive information specifying a limit on how much disk space is available for physical indexes.

- [c32] 32. The system of claim 31, wherein the index consultant takes into account the limit on disk space available for physical indexes.
- [c33] 33. The system of claim 31, wherein the index consultant removes some of the physical indexes from consideration, when sufficient disk space is unavailable.
- [c34] 34. The system of claim 33, wherein the physical indexes removed from consideration are ones having less favorable cost benefits for the captured workload.
- [c35] 35. The system of claim 33, wherein the physical indexes removed from consideration comprise at least 20 percent of bottom performing indexes considered for recommendation.
- [c36] 36. The system of claim 24, wherein the index consultant allows user input specifying whether certain types of indexes should be considered at all.
- [c37] 37. The system of claim 24, wherein the index consultant searches for relevant indexes that will help the system's optimizer use sargable predicates for partial index scans.
- [c38] 38. The system of claim 37, wherein the optimizer creates virtual indexes without specifying ordering of

columns used in sargable equality predicates.

- [c39] 39. The system of claim 24, wherein the index consultant searches for relevant indexes that will help provide useful interesting (order or grouping) properties.
- [c40] 40. The system of claim 39, wherein columns of indexes created may reflect order-independent "don't care" columns that satisfy some interesting ordering wish lists of the system's optimizer.
- [c41] 41. The system of claim 24, wherein the index consultant attempts to collapse some of the virtual indexes together, if feasible for the workload.
- [c42] 42. The system of claim 41, wherein the index consultant attempts to identify that columns of one index are a superset of the columns of another index, and that both indexes may be combined into a single index that is feasible for the workload.
- [c43] 43. The system of claim 24, wherein operation of the index consultant may be polled during operation to ensure that the system is working with accurate cost information.